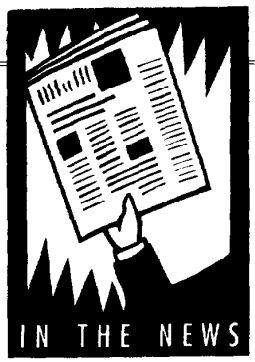

How software affects the world. How the world affects software.

David Sims

INTERNET PROVIDERS FACE LIFE WITHOUT NSFNET

WHEN THE NSFNET WAS TURNED OFF at midnight on April 30, the Internet passed another milestone in its evolution from a government-funded experiment to a commercial endeavor.

Since the US National Science Foundation launched it in 1985, NSFnet has served as a national data "backbone," offering free access to the Internet for US research and educational institutions and supporting the development of regional networks to bring electronic traffic to the backbone from outlying areas. The NSF also helped organize the regional networks, supported protocols, and regulated access agreements between providers.

THE AGENCY MOVES ON TO HIGH-SPEED NETWORK RESEARCH.

But in 1992, the NSF began warning players in the system that it would soon withdraw from "commodity-level network service" because these services could be provided commercially.

"It doesn't make any sense for the government to compete with private businesses that can provide this service," said Beth Gaston, a spokesperson for the NSF. "We see our community as the research and education community, and our research dollars could be better spent on developing new technology."

Along those lines, the NSF announced a five-year cooperative research-and-development agreement with MCI to build a 155-Mbps very-high-speed network to serve its supercomputing centers. The project is described in the box below.

NEW 'NET. NSFnet served as a model for companies developing their own data networks, and these service providers — such as MCI, Sprint, UUNet, and Advanced Network and Services — are now taking over backbone services. Regional network providers, such as Nysernet in New York and Midnet in the Midwest, form peer-to-peer relationships with each other, with other commercial and government networks, and with the commercial backbone providers, sometimes called network service providers.

As Figure 1 shows, the network service providers exchange data at network-access points in San Francisco, Chicago, New York, and Washington, DC. There are plans to add as many as six more NAPs in the future. Data exchange at these NAPs is facilitated by routing

NEXT GENERATION 'NET

On April 24, the NSF and MCI signed a $50-million cooperative research-and-development agreement to develop the very high-speed backbone network service (vBNS) — the next generation of Internet transport that will carry TCP/IP traffic at 155 Mbps (the OC-3 fiber-optic standard).

The vBNS will link the NSF's supercomputing centers: the National Center for Supercomputing Applications at Urbana, Illinois, the San Diego Supercomputing Center, the Pittsburgh Supercomputing Center, the Cornell Theory Center in Ithaca, New York, and the National Center for Atmospheric Research in Boulder, Colorado.

The agreement calls for MCI to boost the vBNS's speed to 622 Mbps in 1996 (OC-12 fiber-optic standard). The NSFnet's fastest transmission speed is over a T3 line, 45 Mbps. The vBNS network will connect the IP through an ATM switching matrix and run this combination on a synchronous optical network (Sonet). The NSF wants to develop the technology to allow 2.2 Gbps transmissions over the network by the end of the five-year agreement. MCI says its Sonet network can operate in a dedicated environment at 2.5 Gbps, but it will take some time — and improvements in routing — before it can bring that performance to the Internet.

"Think of it like an airline flight," said MCI spokesperson Pam Small, "You're really moving when you're in the air; it's landing and changing planes that slows you down."

The NSF continues to authorize use of the research backbone to supercomputer users. Beth Gaston said the agency expects vBNS technology to trickle down to commercial users, as did previous NSFnet technologies, such as T1 (1.5 Mbps) and T3 connectivity.

Editor:
Angela Burgess
IEEE Software
10662 Los Vaqueros Cir.
PO Box 3014
Los Alamitos, CA 90720
aburgess@computer.org

arbiters, which help solve conflicts between network entities. The NSF previously handled this arbitration. It recently signed five-year cooperative agreements worth $10 million each with Merit Network and the University of Southern California's Information Sciences Institute to act as routing arbiters.

LOST IN CYBERSPACE. The transition has not been entirely smooth. Internet traffic patterns — which were relatively simple under NSFnet — are now complex. For example, CERFnet, based in San Diego, California, has identified at least six carriers for its Internet traffic. Pushpendra Mohta, executive director of CERFnet, says the change makes the exterior-routing configuration very complicated.

"Not only are we peering with multiple providers, we are also peering with large national providers at multiple places," Mohta says. Tracking misrouted packets is also more challenging: "It used to be that if your packets did not make it to their destination, you could call the NSFnet Network Operation Center and they could locate the problem. Now this traffic could be routed through multiple exit points. This makes the task of troubleshooting very complicated."

Multiple carriers imply multiple agreements, some made hastily and without clear delineation of the services that would result. Daniel Birchall is a consultant with the International Internet Association, a nonprofit organization that works to ensure rural connectivity. In the two months prior to the NSFnet pullout, Birchall heard reports of frequent routing problems, caused by poor or nonexistent agreements between carriers, that resulted in "chaos."

"When the NSF pulled out," Birchall says, "everybody had to go out and make their own arrangements. Various providers at various levels were frantically trying to sign contractual agreements with one another."

Because providers that previously connected through the NSF had to find new commercial connections, their routing addresses may have changed. As with regular postal service, circulating an address change through the Internet takes time. Other problems may have resulted, Birchall said, because providers had not completed arrangements to accept each other's data.

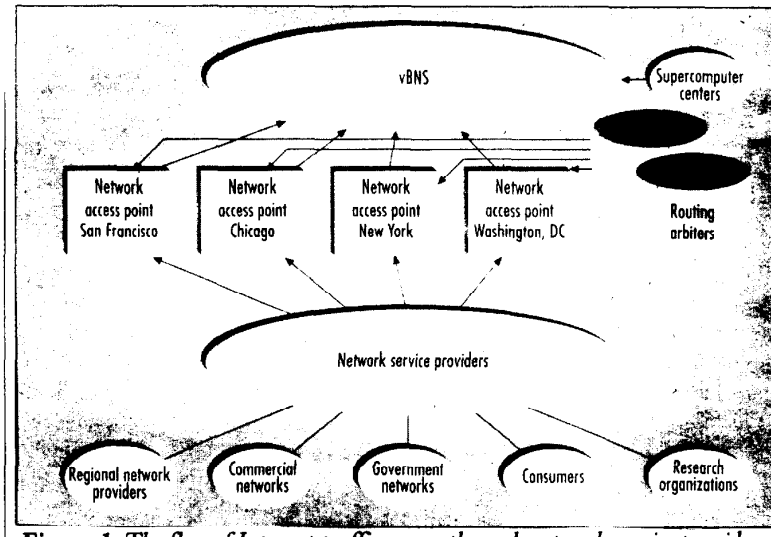

***Figure 1.** The flow of Internet traffic moves through network service providers to network access points and then to the main high-speed network.*

And then there is the matter of personal politics. "I've actually seen situations in which two regional providers wouldn't route each other's data because the person in power at one place didn't like the person in power at the other," Birchall said. He added, however, that since May the bugs appear to be shaking out of the system, and the community could once again become "one big happy family," albeit a family without a head.

The Internet Engineering Task Force continues to define protocol definitions, but beyond basic arbitration and the emerging leadership of various commercial providers, at present no agency oversees the interactions among the system players. The FCC, which is heavily involved in efforts to restructure telecommunications policy as it relates to telephony and television, has little to say on the Internet's structure.

WEANING THE REGIONALS. The NSF has not left without a trace. Rather than eliminating its support of regionals in one blow, the NSF is phasing out its support by 20 percent each year. Also, in an attempt to foster their growth, the NSF has been instructing regionals to behave more like businesses.

Some seem to have taken that advice to heart. Although regionals once respected casual geographic boundaries between one another, a few are now aggressively moving into lucrative territories. BBN Planet, a subsidiary of ARPANet developer Bolt, Beranek and Newman, has acquired three of the largest and most strategically placed regionals, NEARnet of New England, BARRnet of Northern California, and SURAnet of the Southeast. Although these mergers may fuel growth, some in the Internet community are wary of this commercialization.

"I think there's a great unknown as to whether the truly good community things of the Internet will continue," said Martha Stone-Martin, program director at FARnet, the Federation of American Research Networks. Stone-Martin said the members of the NSFnet community are accustomed to an environment in which "you could call up another organization and know the person at the other end of the phone, and solve a particular operational or technical problem . . . based on a common need to make it work."

Even as Internet providers grow in number and size and learn to compete with each other, the need to "make it work" remains. "There's strong support for trying to maintain some level of common goals," Stone Martin said. "If a part of it doesn't work, then it doesn't work."◆

Interview with Watts Humphrey

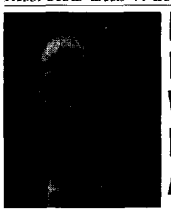

PUSHING PROCESS WITH A PERSONAL APPROACH

WATTS S. HUMPHREY, A RESEARCH scientist at the Software Engineering Institute of Carnegie Mellon University, is no stranger to the readers of *IEEE Software*. The May 1995 issue carried a review of his new book, *A Discipline for Software Engineering* (Addison-Wesley, 1995). In our July 1991 issue, Humphrey reported on his experience at Hughes Aircraft, where he helped improve the development process using the Capability Maturity Model. Recently, Contributing Editor Ware Myers asked him to share his thoughts on how working software engineers might improve their software process.

Q: *You have been working for several decades to improve the process of developing software, at IBM and then at the Software Engineering Institute. At first you emphasized improving the organization. Now your emphasis is on the individual engineer. What led you to the personal software process?*

A: Six years ago at the SEI, I was working on how to apply the Capability Maturity Model to small organizations and projects. At the time, the general view was that the CMM was primarily applicable to large organizations. I decided to experience for myself what a small project is like. I started by defining a personal process as close to CMM level 5 as I could. Then I wrote 62 programs in Object Pascal and C++ — 25,000 lines of code in all. In the course of doing this, I accumulated a lot of data and I learned a great deal. I thought it was a good approach for other people to take.

Q: *What kind of reaction did you get at first?*

A: When I tried to get people to follow these methods, they listened politely, but I had no luck at all. I came closest with a couple of graduate students. They made an attempt, but they just couldn't handle it. It was like teaching somebody to swim at the deep end of the pool.

Q: *How did you get over to the shallow end?*

A: It was evident that people have to get in a step at a time. This led me to the idea of starting with an initial process and then improving it. As I was debating with myself what initial programs people should write, I came up with the idea of having them write simple tools that they could use in a series of programming exercises. That, in turn, led me to the idea of having them use their own data to analyze the effectiveness of the process they were using.

Q: *I realize that you need a 789-page book and a semester-long course to make these steps work, but can you give us some idea of what the personal software process is?*

A: The principles are classic process management: define the work, establish measures, plan the work, measure the work as you do it, analyze the measures to see how to improve, change the process accordingly, and start over.

Q: *Don't engineers learn these principles in school?*

A: Some do, but usually in a manufacturing context. Applying the principles to a knowledge-based process like software development is a new idea. Working engineers have a long history of personal practices that they know work for them. They are especially reluctant to take chances with new techniques in the middle of a project. They may be intellectually impressed with the ideas, and they may agree to experiment with them when they have time — but that time almost never comes.

Q: *How do you get from the principles to the personal software process?*

A: We start with the process the engineers are using. We then add very simple planning and project-management phases that require some simple measurements, such as the time spent in each phase of an exercise and certain data on each defect removed. They have no trouble with this first step. The data they are collecting gives them feedback on their own work.

As part of the course, the engineers then produce periodic reports. Reports 1 and 2 concern line-of-code counting and coding standards. In Report 3, they analyze the defect data accumulated on the first three programming exercises. After the first six programming exercises, they define a process for producing a report and follow this process to produce Report 4. What they learn is that process development is not that hard and that processes apply to lots of things — reports as well as programs.

Q: *All this reporting has given you a lot of data.*

A: Yes. I now have data on more than a hundred engineers and more than a thousand PSP programs. Data from the first 47 people to complete the course — I am still analyzing the rest — shows that they reduced defects by 61.3 percent per thousand lines of code. In test, defects were reduced by 75.5 percent because the students found the defects earlier. Estimating accuracy improved in terms of both size and time. And overall productivity increased by an average of 22 percent.

THE NEW IDEA BEHIND THE PSP IS THE INDIVIDUAL APPLICATION OF MANAGEMENT PRINCIPLES TO SOFTWARE.

But numbers like these are only part of the story. The other part is that the PSP excites those who go through it. They are converted to true believers. It is common for graduates to become disciples — they want to convince their organizations and teammates to use it.

Q: *What does this excitement lead to?*

A: The PSP can help overcome several problems with software development. First, engineers typically don't use known and effective software-development practices. PSP demonstrates to them *with their own data* that these methods can help them do better and faster work. It shows them that they don't need fancy tools, languages, methods, or management support. Second, most engineers use testing as a way to manage product quality; PSP demonstrates to them quantitatively that inspections can

help significantly.

Today, many groups have difficulty defining their processes; they don't have an intuitive understanding of what a process is or how it can help them. PSP shows that defining and using processes is easy and natural. Another common problem is measurement; PSP starts right out making measurements and then uses them.

Finally, most organizations do not follow an orderly or scientific approach to improvement. They have no data to show which problems are most important to them. When they do try new methods, they fail to track them to see if they work. PSP provides ways to measure improvement.

Q: *What role can the individual engineer play in these organizational problems?*

A: The engineer's common view today is that the process is somebody else's job. Engineers recognize that it's all screwed up, but they think that somebody else is responsible for fixing it. Until that somebody fixes it, the engineers just keep plugging away. With the PSP, they learn that a lot of the process is something *they* can fix, with no approvals, support, or tools. This can-do spirit is what we need.

Q: *Do your graduates see it that way?*

A: They sure do. Some of my students have given me feedback. Here are some sample evaluations of PSP by former students:

- "I am very pleased with the convergence of the line-of-code estimates for the final three programs and by the accuracy of both the time and LOC estimate for program 10."
- "By far the most encouraging results from my PSP experiments have been the drastic reduction in my error rate — the injection rate was reduced by fourfold between assignment 1 and assignment 10."
- "Improved effort and source code estimation to within 3 percent of actuals; an 85 percent decrease in post-test defects. . . . My overall performance improved about 15 percent."

Q: *Are these results typical of what engineers can expect after taking the PSP course?*

A: Well, not exactly. Although everyone seems to gain something from the course, their results vary enormously. To me, this course shows how varied software people really are and what a wide range of talents they have. The PSP course will help them to learn more about themselves and where their personal talents lie.

RESULTS FROM THE PSP COURSE VARY WIDELY. ALL DEVELOPERS LEARN MORE ABOUT WHERE THEIR OWN TALENTS LIE.

Q: *Where can people go to learn more?*

A: The SEI helps software organizations introduce the course and also provides courses. I suggest contacting Jim Over (jwo@sei.cmu.edu). A growing number of universities also teach PSP, and Digital Equipment offers several PSP courses. The contact at Digital is John Eikenberry (john.eikenberry@zko.mts.dec.com). Dan Roy, who worked with me at the SEI, also teaches the PSP. He can be reached at (412) 934-0943.

Q: *It is evident that the PSP is already having an effect on a few hundred engineers and their organizations. Assuming that it continues to spread, what broader effect do you see?*

A: Generally speaking, large programs are developed in small parts that are then integrated. When the quality of these small parts is poor, it sharply reduces the quality and productivity of the entire activity. By helping engineers see how to establish and manage high personal quality goals, both organizational productivity and product quality will be significantly improved.

THE ENGINEER'S COMMON VIEW TODAY IS THAT FIXING THE SCREWED UP PROCESS IS SOMEBODY ELSE'S JOB.

A lot of software is developed in teams. To be fully effective, teams must be able to count on their members to perform as they commit. Currently, most software engineers do not know how to plan, estimate, and schedule their work. This inability makes software engineering essentially a solo activity, because the engineers can't count on other team members to deliver a working product when they say they will. When commitment is improved at the individual level, it will be greatly advanced at the team and project levels, too.

Q: *After all the doom and gloom about the software crisis, do you see a brighter future?*

A: There are problems in software development today, but I don't want to leave the impression that it is all the fault of today's software professionals. I don't think we provide the proper engineering disciplines during their formal education. As a result, they struggle with the development of enormously complex systems largely by intuition. The fact that they get these systems to work at all is a remarkable achievement! Unfortunately, society needs much more. Remember, the skill of writing programs is beyond most of the world's population. Software people have encountered one of the most demanding challenges known to mankind. And they have done it largely by themselves.

Q: *Agreed, but back to the future.*

A: The path to a better future certainly beckons us. I have been surprised to see how effective the PSP approach has been. By getting engineers to try new methods and gather data on what they do, their own data convinces them that they need to change the way they work. Now I suddenly see engineers doing things in a matter of weeks that they never would have considered before.

The PSP can show engineers how to practice 12 of the 18 key process areas in the CMM. Basically, that's all of the practices in level 4 and 5 and most in level 3. It can help organizations build more balanced and effective teams. It can provide organizations with the data they need to consistently improve their performance. Still, individual engineers have to take the first steps toward their own personal software process. ◆

♦ **Corel pays beta testers.** Corel has begun paying some of its beta testers, who have traditionally been volunteers whose only reward is free use of new software and a peek at what's next. The move is part of efforts to beef up the company's beta-testing program. Corel says it's the first company to pay its beta testers.

— *NewsBytes News Network*

♦ **AAS endures.** The US Federal Aviation Administration has decided to press on with its troubled Advanced Automation System if the contractor, Loral Federal Systems, fires the project's management team. Loral's $955 million contract — mostly for new air-traffic-control workstations — is its largest contract award since it acquired the former IBM Federal Systems business last year. FAA Deputy Administrator Linda Daschle said Loral is being given a chance to salvage the AAS project, which has already cost $1.4 billion and is expected to reach $5.4 billion. Loral will also get $57 million for an automated system for control towers. If a prototype for the El Paso airport works, the FAA will negotiate with Loral for 69 more systems.

— *NewsBytes News Network*

♦ **PSM Draft released at STC.** The coordination draft of *Practical Software Measurement: A Guide to Objective Program Insight* was released at the US Department of Defense's annual Software Technology Conference in April. The PSM draft is available for public comment prior to approval by the Joint Logistics Commanders, the sponsors of the effort. To get a copy, contact Jack McGarry at mcgarry@ada.npt.nuwc.navy.mil. Proceedings of the STC are published on a CD-ROM. For more information, contact Dana Dovenbarger at dovenbar@oodis01.hill.af.mil.

♦ **Superhighway standards.** A panel set up by the American National Standards Institute has produced draft descriptions of 16 standards requirements for the infrastructure of the information superhighway. The Information Infrastructure Standards Panel developed the drafts in April and is distributing them to more than 30 standards groups and consortia. The standards address issues like reliability, network-to-network interfacing, security, quality of service, and set-top box interface. R.M. "Chick" Hayden, director of information infrastructure programs at ANSI, said the panel is looking for feedback from other standards groups to see if some requirements are met by existing standards. The panel has set up working relationships with comparable groups in Canada, France, Italy, Japan, and Sweden to collaborate on global implementation.

— *NewsBytes News Network*

♦ **Bellcore for sale.** The seven regional Bell operating companies that provide local telephone service across the US plan to sell Bellcore, their shared software, consulting, and research operation. Ameritech, Bell Atlantic, BellSouth, Nynex, Pacific Telesis, SBC Communications, and US West have owned Bellcore since 1984, when it was set up in the wake of the breakup of AT&T to provide common products and services for the Baby Bells. As deregulation allows the RBOCs to move into competitive positions, Bellcore spokesman Ken Branson said, their needs are diverging. "We are what they have in common, and they have less and less in common every day," he said. In the past two years, Bellcore has nearly doubled its sales to companies other than its RBOC owners; outside sales make up nearly 17 percent of the firm's revenue. Bellcore had revenues of more than $1 billion in 1994, and employs about 6,000 people. Branson noted that although the company is primarily thought of as a research organization, only about 10 percent of its organization is basic research; the rest is applied research in software development and consulting services.

— *NewsBytes News Network*

♦ **Standards for WWW searches.** The architect of Carnegie Mellon University's Lycos World Wide Web search engine and architects of other search engines, such as Aliweb, Web Arts, RBSE, and Web Crawler, have proposed forming a body to draft search-engine standards. Michael Mauldin said he and the other search designers met at the Braustubl pub during a Web conference in Germany in April and came up with the idea to standardize the increasingly important process of Web searching. Lycos' popular Web search site is moving its 35 Gbytes of text from the CMU network to its own URL, http://www.lycos.com. Mauldin said that although it will continue to offer some free services, it may become more of a commercial online-information repository, a role it will play for Microsoft Network. Another popular Web catalog, Yahoo, has grown in one year from a part-time hobby for two Stanford University students to Yahoo, Inc., receiving investment funds from Sequoia Capital. The Yahoo guide at http://www.yahoo.com draws 20,000 users a day who browse through 10 million pages a week. Developers Jerry Yang and David Filo say they intend to use the backing from Sequoia to beef up the system's infrastructure. They hope to keep the site free to users by eventually incorporating advertising.

— *NewsBytes News Network*

♦ **Bombing stirs Internet debate.** In the wake of the bombing in Oklahoma City, FBI officials have renewed their campaign for controls on public encryption, and some members of Congress have called for restrictions on Internet speech. In testimony to the Senate Judiciary Committee, FBI Director Louis Freeh warned that terrorists could use encryption technology such as Pretty Good Privacy to communicate on the Internet. Freeh asked for more access to telecommunications carriers "for counterterrorism and counterintelligence investigations." He appeared to win a convert in Sen. Edward Kennedy (D-Mass.), usually considered a free-speech defender. Kennedy held up what he called a "how-to manual for the terrorist" that a staff member had downloaded from the Internet, and said, "We know the advantages of the information highway, but there's the darker side we need to be concerned about." Civil libertarians are concerned that Congress could impose limits on the anarchic, free-spirited Internet. Writing in the *New York Times*, Electronic Frontier Foundation board member Denise Caruso quoted Benjamin Franklin: "They who can give up essential liberty to obtain a little temporary safety deserve neither liberty nor safety."

— *NewsBytes News Network*

♦ **A PC in every pot.** *Computer Industry Almanac* reports that the US now has one computer for every three citizens, up threefold from 1985. The ratio is expected to grow to four computers for every 10 people by the end of 1995. Most of those machines are PCs, but the count includes mainframes and workstations. The US computer industry sells 1.5 million computers each month, the *Almanac* reported. The worldwide total is about one-tenth of the US's, with 35 machines for every thousand people.

— *NewsBytes News Network*